Empirical study of PROXTONE and PROXTONE$^+$ for Fast Learning of Large Scale Sparse Models

Ziqiang Shi and **Rujie Liu**[1]

arXiv:1604.05024v1 [cs.LG] 18 Apr 2016

Abstract. PROXTONE is a novel and fast method for optimization of large scale non-smooth convex problem [18]. In this work, we try to use PROXTONE method in solving large scale *non-smooth non-convex* problems, for example training of sparse deep neural network (sparse DNN) or sparse convolutional neural network (sparse CNN) for embedded or mobile device. PROXTONE converges much faster than first order methods, while first order method is easy in deriving and controlling the sparseness of the solutions. Thus in some applications, in order to train sparse models fast, we propose to combine the merits of both methods, that is we use PROXTONE in the first several epochs to reach the neighborhood of an optimal solution, and then use the first order method to explore the possibility of sparsity in the following training. We call such method PROXTONE plus (PROXTONE$^+$). Both PROXTONE and PROXTONE$^+$ are tested in our experiments, and which demonstrate both methods improved convergence speed twice as fast at least on diverse sparse model learning problems, and at the same time reduce the size to 0.5% for DNN models. The source of all the algorithms is available upon request.

1 INTRODUCTION

Benefited from the advances in deep learning and big data, the accuracy has been dramatically improved on difficult pattern recognition problems in vision and speech [8, 5]. But currently there are two urgent problems need to solve for real life, especially internet applications of deep learning: the first one is that it always took a very long time to adjust the structures and parameters to obtain a satisfactory deep model; and the second one is how to program the always really big deep network on embedded devices or mobile devices. Thus fast learning of sparse regularized models, for example, such as L1 regularized logistic regression, L1 regularized deep neural network (sparse DNN) or L1 regularized convolutional neural network (sparse CNN) becomes very important.

In order to solve the problem of learning large scale L1 regularized model:

$$\min_{x\in\mathbb{R}^p} f(x) := \frac{1}{n}\sum_{i=1}^{n} g_i(x) + \lambda_2\|x\|_1, \tag{1}$$

researchers have proposed the standard and popular *proximal stochastic gradient descent* methods (ProxSGD), whose main appealing is that they have an iteration cost which is independent of n, making them suited for modern problems where n may be very large. The basic ProxSGD method for optimizing (1), uses iterations of the form

$$x^k = \mathcal{S}_{\alpha_k h}[x^{k-1} - \alpha_k \nabla g_{i_k}(x^{k-1})], \tag{2}$$

where $\mathcal{S}_\varepsilon[\cdot]$ is the soft-thresholding operator:

$$\mathcal{S}_\varepsilon[x] \doteq \begin{cases} x-\varepsilon, & \text{if } x > \varepsilon \\ x+\varepsilon, & \text{if } x < -\varepsilon, \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

and at each iteration an index i_k is sampled uniformly from the set $\{1, ..., n\}$. The randomly chosen gradient $\nabla g_{i_k}(x_{k-1})$ yields an unbiased estimate of the true gradient $\nabla g(x_{k-1})$ and one can show under standard assumptions that, for a suitably chosen decreasing step-size sequence $\{\alpha_k\}$, the ProxSGD iterations have an expected sub-optimality for convex objectives of [2]

$$\mathbb{E}[f(x^k)] - f(x^*) = O(\frac{1}{\sqrt{k}})$$

and an expected sub-optimality for strongly-convex objectives of

$$\mathbb{E}[f(x^k)] - f(x^*) = O(\frac{1}{k}).$$

In these rates, the expectations are taken with respect to the selection of the i_k variables.

Thus at least in theory, in fact also in practice it is showed that ProxSGD is very slow in solving the problem 1. While in real life applications, we need to learn and adjust fast in order to obtain a usable model quickly. This requirement results in a large variety of approaches available to accelerate the convergence of ProxSGD methods, and a full review of this immense literature would be outside the scope of this work. Several recent work considered various special or general cases of (1), and developed algorithms that enjoy the linear convergence rate, such as ProxSDCA [16], MISO [11], SAG [15], ProxSVRG [20], SFO [19], ProxN [10], and PROXTONE [18]. All these methods converge with an exponential rate in the value of the objective function, except that the ProxN achieves superlinear rates of convergence for the *solution*, however it is a batch mode method. Shalev-Shwartz and Zhang's ProxSDCA [17, 16] considered the case where the component functions have the form $g_i(x) = \phi_i(a_i^T x)$ and the Fenchel conjugate functions of ϕ_i can be computed efficiently. Schimidt et al.'s SAG [15] and Jascha et al.'s SFO [19] considered the case where $\lambda_2 \equiv 0$.

In order to solve the problem (1) with linear convergent rate, we has proposed a novel and fast method called **prox**imal s**to**chastic **N**ewton-type gradient descent (PROXTONE) [18]. Compared to previous methods, the PROXTONE like other typical quasi-Newton techniques, requires no adjustment of hyperparameters. And at the

[1] Fujitsu Research & Development Center, China, email: {shiziqiang, rjliu}@cn.fujitsu.com

same time, the PROXTONE method has the low iteration cost as that of ProxSGD methods, but achieves the following convergence rates according to the two theorems in [18]

$$\mathbb{E}[f(x^k)] - f^* = O(\mu^k \|x^* - x^0\|^2). \quad (4)$$

When some additional conditions are satisfied, for example $\nabla^2 g_i$ are Lipschitz continuous and so on, then PROXTONE converges exponentially to $x^\star$ in expectation

$$\mathbb{E}[\|x^{k+1} - x^\star\|] = O(\eta^k \|x^* - x^0\|^2).$$

For details and proofs, please refer to our previous theory work [18].

The PROXTONE iterations take the form $x^{k+1} \leftarrow x^k + t_k \Delta x^k$, where Δx^k is obtained by

$$\Delta x^k \leftarrow \arg\min_d d^T(\nabla_k + H_k x^k) + \frac{1}{2} d^T H_k d + \lambda_2 \|x^k + d\|_1, \quad (5)$$

here $\nabla_k = \frac{1}{n}\sum_{i=1}^n \nabla_k^i$, $H_k = \frac{1}{n}\sum_{i=1}^n H_k^i$, and at each iteration a random index j and corresponding H_{k+1}^j is selected, then we set

$$\nabla_{k+1}^i = \begin{cases} \nabla g_i(x^{k+1}) - H_{k+1}^i x^{k+1} & \text{if } i = j, \\ \nabla_{k+1}^i & \text{otherwise.} \end{cases}$$

and $H_{k+1}^i \leftarrow H_k^i$ $(i \neq j)$.

In this work, we try to use the second order method PROXTONE to promote the training of sparse deep models. Compared to conventional methods, PROXTONE can make full use of the gradients, thus needs less gradients (epochs) to achieve same performance, which means converges much fast in the number of epochs. But for each gradient, PROXTONE needs to update the hessian, to construct the low-dimensional space, and solve some kind of lasso subproblem, thus needs much more CPU time against first order methods. That means, finally PROXTONE may converges slow in time than first order methods. In order to overcome this problem, in each iteration, we performance less iterations in solving the subproblems, which means we are satisfied with less exact steepest search directions. This approximation accelerate the convergence of PROXTONE, but result in less sparsity in weights of deep neural networks.

During the empirical study, we found that in some situations, for example training of fully connected DNN, fast approximated PROXTONE cannot fully explore the possibility of sparseness in weights. While first order method is easy in deriving and accumulating the sparseness in each iteration by soft threshold operators, thus we propose to combine first order method with PROXTONE in training DNN. We call such kind of methods PROXTONE$^+$. Experiments show that PROXTONE and PROXTONE$^+$ are suitable for training different kind of neural networks, for example PROXTONE is much suitable for sparse CNN, since whose almost all weights are of shared type, while PROXTONE$^+$ is much more suitable for training of sparse DNN. Finally, the optimizer and the code (matlab and python) reproduce the figures in this work is available upon request.

We now outline the rest of this study. Section 2 presents the main PROXTONE algorithm for L1 regularized model learning, and states choice and details in the implementation. Section 3 describe the PROXTONE$^+$ method. We report some experimental results in Section 4, and provide concluding remarks in Section 5.

2 ALGORITHM

Our goal is to use the PROXTONE for sparse regularized model learning. In general, we always separate the n training samples into M, for example several hundred mini-batches, but in order for the simplicity of notations and description, we did not distinguish between n and M. That is in the following algorithms, n means the number of mini-batches, which should be keep in mind. In this section, we first describe the general procedure by which we optimize the parameter x. We then describe the procedure of the BFGS [13] method by which the online Hessian approximation is maintained for each batch or subfunction. This followed by a description of solving the subproblem in PROXTONE.

2.1 PROXTONE

In each iteration, general PROXTONE uses a L1 regularized piecewise quadratic function to approximate the target loss function for the deep model in a local area around the current point x^{k+1}, and the solution of the regularized quadratic model is used to be the new point. The component function $g_{i_k}(x)$ is sampled randomly, and then the gradient and the approximation of the hessian is used to update the the regularized quadratic model. The procedure is summarized in the Algorithm 1.

Algorithm 1 PROXTONE for L1 regularized model learning

Input: start point $x^0 \in \text{dom } f$; for $i \in \{1, 2, .., n\}$, let $H_{-1}^i = H_0^i$ be a positive definite approximation to the Hessian of $g_i(x)$ at x^0, $\nabla_{-1}^i = \nabla_0^i = \nabla g_i(x^0) - H_0^i x^0$, and let $g_i^0(x) = g_i(x^0) + (x - x^0)^T \nabla g_i(x^0) + \frac{1}{2}(x - x^0)^T H_0^i (x - x^0)$; $G^0(x) = \frac{1}{n}\sum_{i=1}^n g_i^0(x)$; $y \in \mathbb{R}^{p*\text{MAX_HISTORY}*n}$, the history of gradient changes for all $i \in \{1, 2, .., n\}$; $last_x \in \mathbb{R}^{p*n}$ and $last_df \in \mathbb{R}^{p*n}$ holds the last position and the last gradient for all the objective functions; MAX_HISTORY; $s \in \mathbb{R}^{p*\text{MAX_HISTORY}*n}$, the history of x or position changes for all $i \in \{1, 2, .., n\}$.

1: **repeat**

2: Solve the subproblem (it is indeed the well known lasso problem) for new approximation of the solution:

$$x^{k+1} \leftarrow \arg\min_x \big[G^k(x) + \lambda_2 \|x\|_1\big]. \quad (6)$$

3: Sample i_k from $\{1, 2, .., n\}$, update the history of position and gradient differences for the mini-batch i_k:

$$s(:, 2 : \text{MAX_HISTORY}, i_k) = s(:, 1 : \text{MAX_HISTORY} - 1, i_k)$$
$$s(:, 1, i_k) = x - last_x(:, i_k)$$
$$y(:, 2 : \text{MAX_HISTORY}, i_k) = y(:, 1 : \text{MAX_HISTORY} - 1, i_k)$$
$$y(:, 1, i_k) = \nabla g_{i_k}(x^{k+1}) - last_df(:, i_k)$$

4: Update the Hessian approximation $H_{k+1}^{i_k}$ for the mini-batch i_k (described in detail in Algorithm 2);

5: Update the quadratic models or surrogate functions:

$$g_{i_k}^{k+1}(x) = g_{i_k}(x^{k+1}) + (x - x^{k+1})^T \nabla g_{i_k}(x^{k+1}) + \frac{1}{2}(x - x^{k+1})^T H_{k+1}^{i_k}(x - x^{k+1}), \quad (7)$$

while leaving all other $g_i^{k+1}(x)$ unchanged: $g_i^{k+1}(x) \leftarrow g_i^k(x)$ $(i \neq j)$; and $G^{k+1}(x) = \frac{1}{n}\sum_{i=1}^n g_i^{k+1}(x)$.

6: **until** stopping conditions are satisfied.

Output: x^k.

In deep learning, the dimensionality of x is always large. As a result, the memory and computational cost of working directly with

the matrices in Algorithm 1 is prohibitive, as is the cost of storing the history terms and required by BFGS. Thus we employ the idea from [19], that is we construct a shared low dimensional subspace which makes the algorithm tractable in terms of computational overhead and memory for large problems. x and the gradients are mapped into a limited sized shared adaptive low-dimensional space, which is expanded when meeting a new observation. The Hessian, the regularized quadratic model, and further the solution are updated in this low-dimensional space. Finally then solution is projected back to the original space to become the real optimal points. This mapping or projection is comprised of a dense matrix, thus the sparse solution in low-dimensional space may result in non-sparse solution in original space. This problem will be discussed and solved in Section 3.

2.2 Hessian approximation

Arguably, the most important feature of this method is the regularized quadratic model, which incorporates second order information in the form of a positive definite matrix $H_k^{i_k}$. This is key because, at each iteration, the user has complete freedom over the choice of $H_k^{i_k}$. A few suggestions for the choice of H_k^i include: the simplest option is $H_k^i = I$ that no second order information is employed; $H_k^{i_k} = \nabla^2 g_i(x_k)$ provides the most accurate second order information, but it is (potentially) much more computationally expensive to work with; in order to do a tradeoff between accuracy and complexity, the most popular formulae for updating the Hessian approximation is the BFGS formula, which is defined by

$$B_k = B_{k-1} - \frac{B_{k-1}s_{k-1}s_{k-1}^T B_{k-1}}{s_{k-1}^T B_{k-1} s_{k-1}} + \frac{y_{k-1}y_{k-1}^T}{y_{k-1}^T s_{k-1}}, \qquad (8)$$

where

$$y_{k-1} = \nabla f(x_k) - \nabla f(x_{k-1}), \quad s_{k-1} = x_k - x_{k-1}.$$

We store a certain number (say, MAX_HISTORY) of the vector pairs $\{s_k, y_k\}$ used in the above formulas. After the new iteration is computed, the oldest vector pair in the set of pairs $\{s_i, y_i\}$ is replaced by the new pair $\{s_k, y_k\}$ obtained from the above step. In this way, the set of vector pairs includes curvature information from the MAX_HISTORY most recent iterations. This is indeed the famous limited-memory BFGS algorithm, which can be stated formally as the following Algorithm 2.

Algorithm 2 Update the Hessian approximation for the mini-batch i_k

Input: MAX_HISTORY = 20, $s \in \mathbb{R}^{p*\text{MAX_HISTORY}*n}$, the history of x or position changes for all $i \in \{1, 2, .., n\}$, $y \in \mathbb{R}^{p*\text{MAX_HISTORY}*n}$, the history of gradient changes for all $i \in \{1, 2, .., n\}$, $H_{k+1}^{i_k} = I$.
1: **for** $j = 1 :$ MAX_ITER
2: $tmpy = y(:, \text{MAX_HISTORY} + 1 - j, i_k)$;
3: $tmps = s(:, \text{MAX_HISTORY} + 1 - j, i_k)$;
4: **if** $tmpy^T * tmps > 0$
5: $H_{k+1}^{i_k} = H_{k+1}^{i_k} - (H_{k+1}^{i_k} * (tmps * tmps^T) * H_{k+1}^{i_k})/(tmps^T * H_{k+1}^{i_k} * tmps) + (tmpy * tmpy^T)/(tmpy^T * tmps)$;
6: **end if**
7: **end for**
Output: $H_{k+1}^{i_k}$.

After the obtaining of $H_{k+1}^{i_k}$, then we can update the local regularized quadratic model (the subproblem), which can be solved by a proximal algorithm.

2.3 The subproblem

The subproblem (10) is a lasso problem, which can be effectively and accurately solved by the proximal algorithms [14]. It is summarized in Algorithm 3.

$$x^{k+1} \leftarrow \arg\min_x \left[G^k(x) + \lambda_2 \|x\|_1\right] = \arg\min_x F^k(x). \qquad (9)$$

Algorithm 3 Solving subproblem (10) based on proximal algorithms

Input: start point $x_0 = x^k$, MAX_ITER = 100, ABSTOL = 1E-5, λ = 1, β = 0.5.
1: **for** i=1:MAX_ITER
2: grad_x = $\nabla G^k(x_{i-1})$,
3: **while** 1
4: $z = \mathcal{S}_{\lambda*\lambda_2}[x_{i-1} - \text{lambda} * \text{grad_x}]$,
5: **if** $G^k(z) <= G^k(x_{i-1}) + \text{grad_x}' * (z - x_{i-1}) + (1/(2 * \lambda)) * \|z - x_{i-1}\|^2$
6: break;
7: **end if**
8: $\lambda = \beta * \lambda$;
9: **end while**
10: $x_i = z$;
11: **if** $i > 1$ && $|F^k(x_i) - F^k(x_{i-1})| <$ ABSTOL
12: break;
13: **end if**
14: **end for**
Output: $x^{k+1} = x_i$.

That means for each gradient, we need to use several iterations of computing approximated Hessian to forming a lasso problem, which also needs several iterations to solve. Thus typically PROXTONE needs much more time for each iteration than that of first order method. That means although PROXTONE is much fast than other other methods in the number of gradients or epochs, but may be slower in time. In the following section, we will try to solve this problem.

3 The PROXTONE$^+$

Compared to conventional method, PROXTONE can achieve the same performance with less gradients, that is in less epochs. But since it always needs much more computation than first order method for each iteration, thus always PROXTONE converges slowly than first order methods in physic time. In order to speed up the PROXTONE, we try to not solve the lasso problem so exactly, that is we always set 'MAX_ITER = 1' in the Algorithm 3. This result in inexact solution in each iteration of PROXTONE, but also result in much faster convergence speed. This speed up cause new problems, that is we cannot control the sparseness of the solution. In order to overcome this problem, we try to combine PROXTONE with first order method, that is in the first stage, we use PROXTONE to reach the nearby of the optimal, and then comes to the second stage, we use ProxSAG to further explore the possibility of sparseness of the solution. The rough idea result in the following PROXTONE$^+$ algorithms.

4 EXPERIMENTAL RESULTS

We compared our optimization technique to several competing optimization techniques for several objective functions. The results are

Algorithm 4 PROXTONE$^+$ for L1 regularized model learning

Input: start point $x^0 \in \text{dom } f$; for $i \in \{1, 2, .., n\}$, let $H^i_{-1} = H^i_0$ be a positive definite approximation to the Hessian of $g_i(x)$ at x^0, $\nabla^i_{-1} = \nabla^i_0 = \nabla g_i(x^0) - H^i_0 x^0$, and let $g^0_i(x) = g_i(x^0) + (x - x^0)^T \nabla g_i(x^0) + \frac{1}{2}(x - x^0)^T H^i_0 (x - x^0)$; and $G^0(x) = \frac{1}{n}\sum_{i=1}^n g^0_i(x)$; N, the number of epochs to perform PROXTONE.

1: **repeat**

2: **if** $k < N$ (use PROXTONE)

3: Solve the lasso subproblem for new approximation of the solution:

$$x^{k+1} \leftarrow \arg\min_x \big[G^k(x) + \lambda_2 \|x\|_1\big]. \tag{10}$$

4: Sample i_k from $\{1, 2, .., n\}$, and update the quadratic models or surrogate functions:

$$g^{k+1}_{i_k}(x) = g_{i_k}(x^{k+1}) + (x - x^{k+1})^T \nabla g_{i_k}(x^{k+1}) + \frac{1}{2}(x - x^{k+1})^T H^i_{k+1}(x - x^{k+1}), \tag{11}$$

while leaving all other $g^{k+1}_i(x)$ unchanged: $g^{k+1}_i(x) \leftarrow g^k_i(x)$ ($i \neq i_k$); and $G^{k+1}(x) = \frac{1}{n}\sum_{i=1}^n g^{k+1}_i(x)$.

5: **else** (use ProxSAG)

6: Sample i_k from $\{1, 2, .., n\}$, and update the gradient $y_{k,i}$ and the average gradient y_{k-1}:

$$y_{k,i} = \begin{cases} \nabla g_i(x_k) & \text{if } i = i_k, \\ y_{k-1,i} & \text{otherwise.} \end{cases} \tag{12}$$

$$y_k = \frac{1}{n}\sum_{i=1}^n y_{k,i} \tag{13}$$

and finally the update of x^{k+1}:

$$x^{k+1} = \mathcal{S}_{\lambda_2/L}[x^k - y_k/L], \tag{14}$$

7: **end if**

8: **until** stopping conditions are satisfied.

Output: x^k.

illustrated in Figures 1, 3, and 4, and the optimization techniques and objectives are described below. For all problems our method outperformed all other techniques in the comparison.

4.1 Sparse regularized logistic regression

In our preliminary study, we use some large scale convex problems to debug our algorithm. Here present the results of some numerical experiments to illustrate the properties of the PROXTONE method. We focus on the sparse regularized logistic regression problem for binary classification: given a set of training examples $(a_1, b_1), \ldots, (a_n, b_n)$ where $a_i \in \mathbb{R}^p$ and $b_i \in \{+1, -1\}$, we find the optimal predictor $x \in \mathbb{R}^p$ by solving

$$\min_{x \in \mathbb{R}^p} \quad \frac{1}{n}\sum_{i=1}^n \log\big(1 + \exp(-b_i a_i^T x)\big) + \lambda_1 \|x\|_2^2 + \lambda_2 \|x\|_1,$$

where λ_1 and λ_2 are two regularization parameters. We set

$$g_i(x) = \log(1 + \exp(-b_i a_i^T x) + \lambda_1 \|x\|_2^2, \qquad h(x) = \lambda_2 \|x\|_1, \tag{15}$$

and

$$\lambda_1 = 1E-4, \qquad \lambda_2 = 1E-4.$$

We used some publicly available data sets. The *protein* data set was obtained from the KDD Cup 2004[2]; the covertype data sets were obtained from the LIBSVM Data[3].

The performance of PROXTONE is compared with some related algorithms:

- ProxSGD (Algorithm 5): We used a constant step size that gave the best performance among all powers of 10;
- ProxSAG (Algorithm 6): This is a proximal version of the SAG method, with the trailing number providing the Lipschitz constant;

Algorithm 5 ProxSGD for L1 regularized model learning

Input: start point $x^0 \in \text{dom } f$.

1: **repeat**

2: Sample i_k from $\{1, 2, .., n\}$,

$$x^{k+1} = \mathcal{S}_{\eta\lambda_2}[x^k - \eta \nabla g_{i_k}(x^k)], \tag{16}$$

3: **until** stopping conditions are satisfied.

Output: x^k.

The results of the different methods are plotted for the first 100 and 500 effective passes for protein and covertype respectively through the data in Figure 1. Here we test PROXTONE with two kinds of Hessian, the first is with diagonal Hessian with constant diagonal elements, and the Hessian of the other kind is updated by Algorithm 2. The iterations of PROXTONE seem to achieve the best of all.

4.2 Sparse deep learning

Two kinds of widely used typical deep learning models, which are sparse DNN and CNN, are used to test our method.

First we trained a deep neural network to classify digits on the MNIST digit recognition benchmark. We used a similar architecture to [6]. The MNIST [9] dataset consists of 28*28 pixel greyscale

[2] http://osmot.cs.cornell.edu/kddcup

[3] http://www.csie.ntu.edu.tw/∼cjlin/libsvmtools/datasets

Algorithm 6 ProxSAG for L1 regularized model learning

Input: start point $x^0 \in \text{dom } f$; let $y_{0,i} = \nabla g_i(x^0)$, and $y_0 = \frac{1}{n}\sum_{i=1}^n y_{0,i}$ be the average gradient.
1: **repeat**
2: Sample i_k from $\{1, 2, .., n\}$, and update the gradient $y_{k,i}$ and the average gradient y_{k-1}:

$$y_{k,i} = \begin{cases} \nabla g_i(x_k) & \text{if } i = i_k, \\ y_{k-1,i} & \text{otherwise.} \end{cases} \tag{17}$$

$$y_k = \frac{1}{n}\sum_{i=1}^n y_{k,i} \tag{18}$$

and finally the update of x^{k+1}:

$$x^{k+1} = \mathcal{S}_{\lambda_2/L}[x^k - y_k/L], \tag{19}$$

4: **until** stopping conditions are satisfied.
Output: x^k.

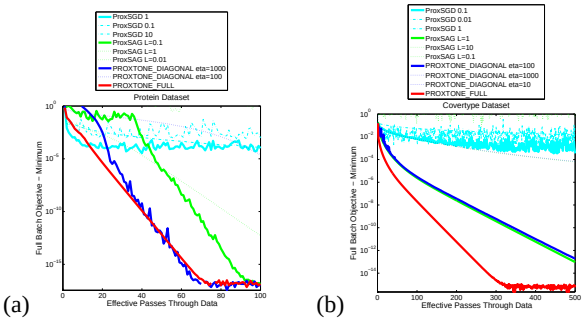

(a) (b)

Figure 1. A comparison of PROXTONE to competing optimization techniques in solving sparse regularized logistic regression for two datasets, (a) is protein; (b) is covertype. The bold lines indicate the best performing hyperparameter for each optimizer.

images of handwritten digits 0-9, with 60,000 training and 10,000 test examples. Our network consisted of: 784 input units, one hidden layer of 1200 units, a second hidden layer of 1200 units, and 10 output units. We ran the experiment using both rectified linear and sigmoidal units. The objective used was the standard softmax regression on the output units. Theano [1] was used to implement the model architecture and compute the gradient.

Second we trained a deep convolutional network on CIFAR-10 using max pooling and rectified linear units. The CIFAR-10 dataset [7] consists of 32*32 color images drawn from 10 classes split into 50,000 train and 10,000 test images. The architecture we used contains two convolutional layers with 48 and 128 units respectively, followed by one fully connected layer of 240 units. This architecture was loosely based on [4]. Pylearn2 [3] and Theano were used to implement the model.

A preliminary experiment is used to choose the hyperparameter of ProxSAG and ProxSGD for sparse DNN and sparse CNN respectively in Figure 2. Then we do detail measurement of time and sparsity for all the methods. The Figure 3 and 4 show that PROXTONE and PROXTONE$^+$ converge nearly twice as fast then the state-of-the-art methods. While for sparsity, PROXTONE$^+$ can reduce the size to about 0.5% for sparse DNN training. Since there are many share weights in CNN, for sparse CNN training, PROXTONE is much more suitable than PROXTONE$^+$, and reduce the size to about 60%.

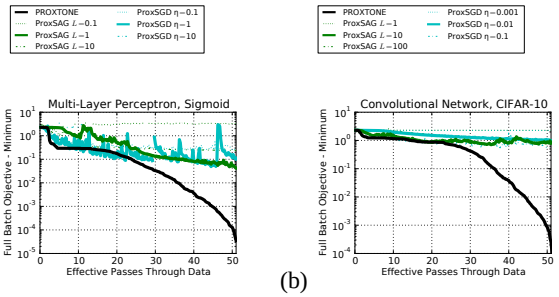

(a) (b)

Figure 2. A comparison of PROXTONE to competing optimization techniques for (a) sparse DNN and (b) sparse CNN. The bold lines indicate the best performing hyperparameter for each optimizer. It can be seen that PROXTONE is free of chosen hyperparameter.

5 CONCLUSION

This paper is to make clear the implementation details of PROXTONE and do the numerical evaluations to nonconvex problems, especially sparse deep learning problems. We show that PROXTONE and PROXTONE$^+$ can make full use of gradients, converges much faster than state-of-the-art first order methods in the number of gradients or epochs. It is also showed the methods converges faster also in time, while reduce the size to 0.5% and 60% for DNN and CNN models respectively. There are some directions that the current study can be extended. Experiments show that ProxSAG method has good performance, thus it would be meaningful to also make clear the theory for the convergence of ProxSAG [15]. Second, combine with randomized block coordinate method [12] for minimizing regularized convex functions with a huge number of varialbes/coordinates. Moreover, due to the trends and needs of big data, we are designing distributed/parallel PROXTONE for real life applications. In a broader context, we believe that the current paper could serve as a

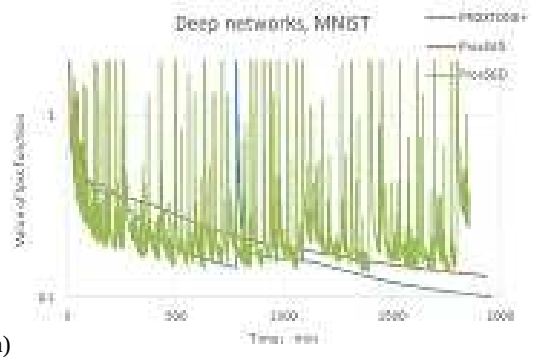

(a)

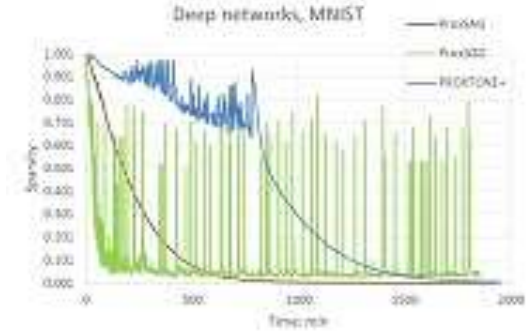

(b)

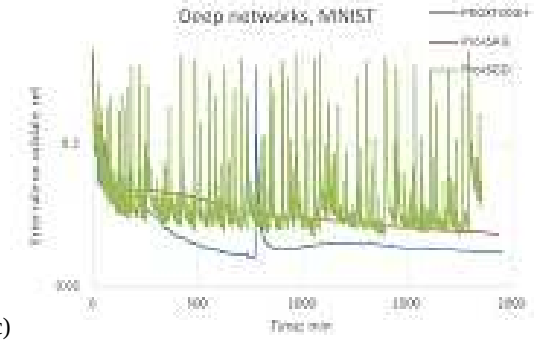

(c)

Figure 3. A comparison of PROXTONE$^+$ to competing optimization techniques for two objective functions. The objective functions shown is a multilayer perceptron with sigmoidal units trained on MNIST digits. (a) value (b) sparsity (c) error.

(a)

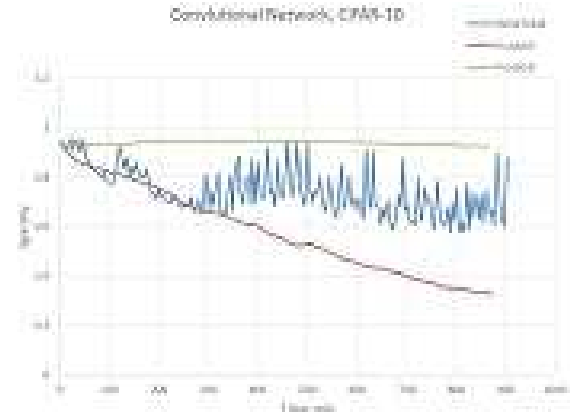

(b)

Figure 4. A comparison of PROXTONE to competing optimization techniques for two objective functions. The objective functions shown is a multilayer convolutional network with rectified linear units trained on CIFAR-10. (a) value (b) sparsity.

basis for examining the method for deep learning on the proximal stochastic methods that employ second order information.

REFERENCES

[1] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio, 'Theano: a cpu and gpu math expression compiler', in *Proceedings of the Python for scientific computing conference (SciPy)*, volume 4, p. 3. Austin, TX, (2010).

[2] Dimitri P Bertsekas, 'Incremental gradient, subgradient, and proximal methods for convex optimization: a survey', *Optimization for Machine Learning*, **2010**, 1–38, (2011).

[3] Ian J Goodfellow, David Warde-Farley, Pascal Lamblin, Vincent Dumoulin, Mehdi Mirza, Razvan Pascanu, James Bergstra, Frédéric Bastien, and Yoshua Bengio, 'Pylearn2: a machine learning research library', *arXiv preprint arXiv:1308.4214*, (2013).

[4] Ian J Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, and Yoshua Bengio, 'Maxout networks', *arXiv preprint arXiv:1302.4389*, (2013).

[5] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al., 'Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups', *Signal Processing Magazine, IEEE*, **29**(6), 82–97, (2012).

[6] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov, 'Improving neural networks by preventing co-adaptation of feature detectors', *arXiv preprint arXiv:1207.0580*, (2012).

[7] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images, 2009.

[8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, 'Imagenet classification with deep convolutional neural networks', in *Advances in neural information processing systems*, pp. 1097–1105, (2012).

[9] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner, 'Gradient-based learning applied to document recognition', *Proceedings of the IEEE*, **86**(11), 2278–2324, (1998).

[10] Jason Lee, Yuekai Sun, and Michael Saunders, 'Proximal newton-type methods for convex optimization', in *Advances in Neural Information Processing Systems*, pp. 836–844, (2012).

[11] Julien Mairal, 'Optimization with first-order surrogate functions', *arXiv preprint arXiv:1305.3120*, (2013).

[12] Yu Nesterov, 'Efficiency of coordinate descent methods on huge-scale optimization problems', *SIAM Journal on Optimization*, **22**(2), 341–362, (2012).

[13] Jorge Nocedal and Stephen Wright, *Numerical optimization*, Springer Science & Business Media, 2006.

[14] Neal Parikh and Stephen Boyd, 'Proximal algorithms', *Foundations and Trends in optimization*, **1**(3), 123–231, (2013).

[15] Mark Schmidt, Nicolas Le Roux, and Francis Bach, 'Minimizing finite sums with the stochastic average gradient', *arXiv preprint arXiv:1309.2388*, (2013).

[16] Shai Shalev-Shwartz and Tong Zhang, 'Proximal stochastic dual coordinate ascent', *arXiv preprint arXiv:1211.2717*, (2012).

[17] Shai Shalev-Shwartz and Tong Zhang, 'Stochastic dual coordinate ascent methods for regularized loss', *The Journal of Machine Learning Research*, **14**(1), 567–599, (2013).

[18] Ziqiang Shi and Rujie Liu, 'Large scale optimization with proximal stochastic newton-type gradient descent', in *Machine Learning and Knowledge Discovery in Databases*, 691–704, Springer, (2015).

[19] Jascha Sohl-Dickstein, Ben Poole, and Surya Ganguli, 'Fast large-scale optimization by unifying stochastic gradient and quasi-newton methods', in *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 604–612, (2014).

[20] Lin Xiao and Tong Zhang, 'A proximal stochastic gradient method with progressive variance reduction', *arXiv preprint arXiv:1403.4699*, (2014).